Data Collection Modules for the PHENIX Experiment

C.Y. Chi, B. Cole, J.L. Nagle, W. Sippach, W.A. Zajc
Nevis Laboratories, Columbia University

Abstract

The data acquisition (DAQ) system for the PHENIX experiment is designed as a pipeline system with simultaneous triggering and readout. The maximum average level-1 (LVL1) trigger rate is 25 KHz. The DAQ system consists of Front-End Modules (FEM's), a level-1 (LVL1) trigger, data collection modules (DCM's), timing systems, slow controllers and an event builder (EVB). The data collection modules have the responsibility of collecting uncompressed LVL1 trigger event fragments from the FEM's. The DCM's provide buffering for up to five LVL1 events. The DCM's also perform zero suppression, error checking, data reformatting and outputting data to the event builder. In addition to the FEM data, the DCM's also receive primitives from LVL1 trigger system. These primitives are used for alignment checking on the FEM data packet. Additional trigger primitives can also be generated together with the FEM data. The DCM is hosted in VME64X crate. VME is used as a means for maintenance and slow control. Data collection within the crate is done through a private data-way in P0 space.

I. Introduction

PHENIX is a one of the two large apparatus experiments to make use of the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory. PHENIX is designed to make measurements on a variety of colliding systems from proton on proton to Au on Au. The occupancy in the detector varies from few tracks total (in p+p interactions) to 10% of all detector channels firing in central Au+Au interactions. The beam crossing clock in the collider will run at 9.43 MHz. At the full beam rate, PHENIX will have a data rate of 2.4 GBytes per second after zero suppression of all detector systems. The relations between colliding system, data rate and data size are shown in Figure 1. The front end electronics design is heavily influence by considerations of proton on proton running where the interaction rate is expected to be approximately 10 MHz (at 10 times upgrade-able luminosity). In Au on Au collisions the interaction rate is 13 KHz with a 2.5% estimated detector occupancy. The back end of the system must be able to process this large data volume. The heavy colliding systems present conditions that large scale high energy experiments do not encounter. These considerations influenced our design in having in the DCM's a large amount of processing power and a buffer pipeline system.

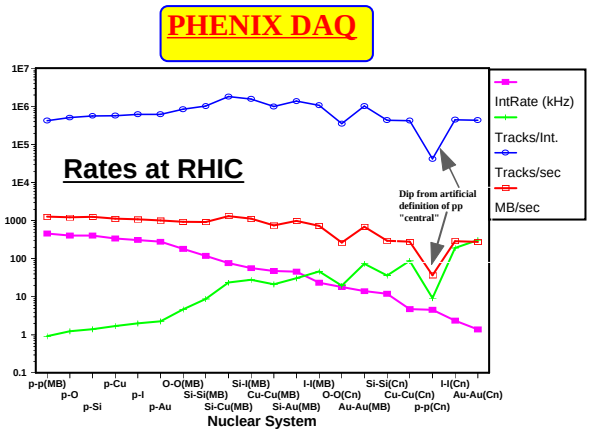

Fig. 1 The relations between collding system, data rate and data sizes.

II. Data Collection Modules

The experiment has roughly 0.3 million channels of electronics. The signal recording methods vary from Y/N discriminators to 40 MHz 5 bit non-linear flash ADC's to Analog Memory Units (AMU's). The front-end electronics are mostly mounted on the detectors which presents limited access constraints. The FEM's continuously sample the signals and store them in either digital or analog memory. After they receive a LVL1 trigger, and after a four microsecond latency, the corresponding event is transferred to the DAQ system. The maximum average LVL1 trigger rate is set to 25 kHz. In order to further increase the trigger efficiency, the system is designed to continuously record the signal while digitizing and transferring the accepted events. Under these conditions the FEM's have to buffer five accepted LVL1 events.

Because there is limited space and significant cooling constraints, we have decided that the FEM's will act as slaves. The FEM's will not perform zero suppression on the accepted LVL1 event. The data passed out of the FEM's will be fixed formatted and transmitted by optical fiber at data rates up to Giga-bits per second. In order to achieve this high bandwidth, flow control on the FEM's will only rely on regulating LVL1 triggers. The FEM's will transfer

data after receiving a LVL1 trigger. Once the DAQ system is full, the DCM's will be responsible for the buffer reset of the FEM's data after stopping LVL1 trigger.

The DCM data flow logical block diagram is shown in Figure 2. There are five major logical units: the compressor, the DSP block, the VME interface, the LVL1 interface and the data output port. Each DCM contains four compressors and DSP modules. There is a fifth DSP used to merge data from four compressors.

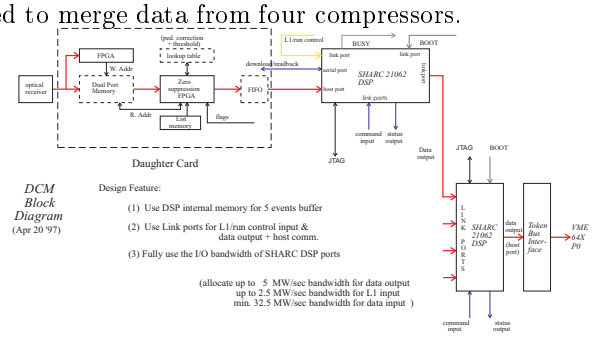

Fig. 2 Logical block diagram of DCM data flow.

Unlike the event builder and higher level triggers, the DCM is viewed as a module attached to a specific FEM. The design has to able to handle the highest possible future data rate. The decision was made to use an (field programmable gate array) FPGA as the computer engine for building different types of compressors. The DSP's are used to handle data formatting, error detection and re-normalization after zero suppression.

A. Compressor Ports

The DCM's perform zero suppression by comparing data with pre-set threshold values. During the normal running conditions, only data above the threshold will be kept. The data transfer rate between FEM and DCM is two or four times the beam crossing clock, i.e. 18.8 or 39.6 MHz. The data words are either 16 bits or 20 bits wide. This data transfer rate is faster than most of commercial DSP's can handle. However, it is still within the FPGA technology ability. There four different types of zero suppression algorithms. The compressors are built as daughter cards to achieve uniformity of the remainder of the DCM board and easy upgrade. AMP 0.8 mm free height 100 pin connectors are used to connect the daughter cards to the main board. The analog memory unit (AMU) type readout will be discussed below.

The AMU's record input signals in an array of capacitors. The array serves both for a LVL1 trigger delay and for buffering of accepted LVL1 data while waiting for digitization. After receiving a LVL1 trigger, the charge on corresponding capacitor is digitized. Different capacitors in the AMU array will have slightly different pedestals. The PHENIX AMU system has 64 capacitor cells in each array. In the DCM, corrections need to applied before threshold values can be compared. The capacitor cell numbers are part of FEM data packet.

The compressor port in the DCM consists of optical receiver, HP GLINK receiver, dual-port memory and compressor FPGA. One list memory and pedestal table are also built into the compressor daughter card. Data is stored in the dual-port memory after being received from optical link. The dual-port memory has 4k and is divided into 2 pages, one for the incoming data while previous event, that stored in the other page, is being processed. An 8K by 16 bits list memory is used by the compressor FPGA to re-group data in the dual-port memory. The lowest 12 bits of the 16 bits data field serve as the dual memory address. The upper 4 bits are used as a label to classify the data. The re-ordering feature is important in order to provide de-coupling of the FEM's and DCM's on the exact data order. The label in the list memory merged with the data from the dual-port memory classified the data into the compressor FPGA as well as DSP.

The compressor FPGA uses channel number, through an incremented counter, and the AMU cell number, contained the data stream, to address the pedestal memory table. The pedestal value will then be removed from then data. The result will compare to the threshold value as a function of channel number. An ALTERA FLEX 10K20 208 pin package is used as the zero suppression engine. The internal memory inside the FPGA is used for the threshold memory table with a maximum size of 512K by 12 bits. The pedestal memory carries 32K by 8 bits. The system will run at 20 MHz.

B. DSP

We chose to use the Analog Device SHARC DSP. SHARC DSP's have the advantage of high speed host ports, six 4 bits wide link ports, and 2 serial ports. The total I/O bandwidth is 40 Megawords/second among all I/O ports. We also make advantage of the large internal memory, both for program and data buffers. The 21062 SHARC DSP is chosen in our case. It has 2 Megabits of dual-port memory. The internal DMA controller and dual-port memory allow the computing unit to operate independently of the I/O. In order to run the compressor port at speed, there is no backflow control. Careful calculations in addition to developing a testing board stations was important to check this feature.

The SHARC DSP's also provide buffering for five FEM data events. This enables us to avoid providing high speed memory on the compressor board. A ring buffer and an automatic DMA chain are used inside the DSP. Busies will be raised once the DSP has less than six uncompressed event buffers. The DSP's have four user defined flags. One of these flags is used as the "busy."

The link ports are 4 bits wide and run at a maximum speed of 40 MHz. This translates to 5 Megawords per second for the case of 32 bits word transfer. The input link port speed is controlled by the sender. The output side is controlled by the DSP clock. Link port 2 is used for receiving LVL1 information. The VME interface uses ports 3 and 4 to read/write data into the DSP. Link port 4 is used, through an FPGA interface, to boot the DSP.

Both the LVL1 data link port and VME input link port are slowed down to reduced the ringing. Link port 5 is used to transfer data between the compressor port DSP and the extra DSP which merges data from the four compressor port DSP's. This limits the average compressor output port to 5 Megawords per second.

C. LVL1 Port

The LVL1 trigger system sends its summary data packet, as well as various clock counters and scalars, to the DCM system every LVL1 trigger. The LVL1 data are 16 bits wide with 1 address tag bit. This data packet is divided into many sub-packets separated by an address word. The address word is recognized by the address tag bit. The upper 8 bits of the address word contain sub-packet identifier information. The lower 8 bits contain the lower 8 bits of the global event counters.

The DCM's filter the LVL1 data word by the upper 8 bits of the address word plus the word position in the packet. The LVL1 data is then received by the crate partition module. It sends data to the DCM through the backplane by a token passing method. The maximum backplane transfer rate is 20 MHz. The DCM's use ALTERA 7K 128 cell FPGA + 64Kx1 SRAM to receive LVL1 data with a short FIFO to back it up. The filter algorithm uses the address word of the packet (8 bits) plus the word counter in the sub-packet (8 bits) as the address word of the 64Kx1 SRAM table. The output of the table, one bit Y/N, will determine whether the data will be passed or not.

The filtered data is passed to a FIFO. The data is then broadcast to the four DSP's on the DCM board through link port 2. One FPGA is used to interface the FIFO and the four DSP's link ports. Figure 3 shows a block diagram of both the DCM-VME and LVL1 interface. One of the LVL1 data sub-packets contains the event counter. By comparing this event counter with FEM event counter contained in the data packet, the DSP's are able to check data alignment within the system. The DSP's use compressed FEM data and LVL1 data in order to re-normalize the FEM data and calculate trigger primitives for the higher level trigger.

D. VME Interface

The DCM's use VME as the path for down loading and slow control/read back. The host interface on the DSP's is used for event data path. The communication from the VME controller to the DSP's has to use the link ports. On each DSP, one port is set up to be the read port and one port is set up to be the write port. One 240 Pin ALTERA FlEX 10K20 chip is used to interface between VMEbus and the DSP link ports. The DCM interface uses VME A24 D16 transfers with a user defined address modifier code. The 24 bits on the address field are decoded into three portions: bit 23 is the broadcast bit, bits 22-18 are the geometric address word, and bits 17-14 are use as a command code. Commands fall into

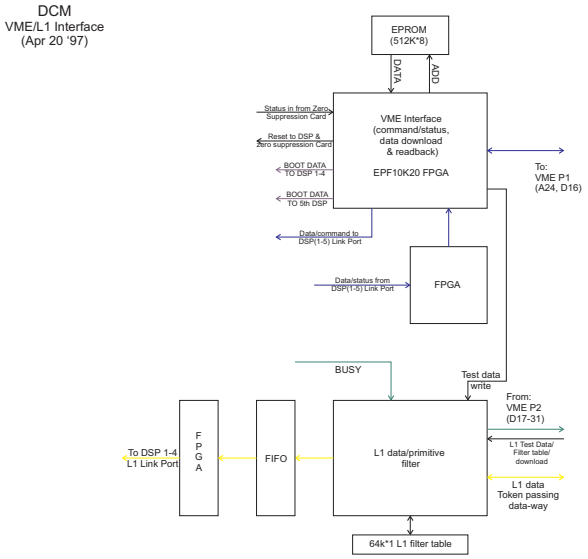

Fig. 3 Block diagram of DCM-VME and LVL1 interfaces.

three class: reading the board status, defining which DSP the interface is connected to, and reading/writing data to a DSP. The SHARC DSP's use 32 bits data and 48 bits program instruction words. It takes two or three 16 bits operations to communicate one DSP word from the controller.

Unlike most of the commercial VME boards, the crate controller and the DSP exchange information through an interface FPGA. The real communication between the DSP's and the VME controller is through the FPGA program. This feature provides flexibility, but it lacks the ability to directly probe the DSP memory. The VME interface FPGA also connects to an EPROM during the power up. A 64kX8 bits EPROM is used. The EPROM contain both DSP 1-4 program and 5th DSP program.

E. Output Port

The fifth DSP collects data from compressor port DSP's (1-4) through link ports. It re-formats the data before sending it out. An ALTERA 7K 192 cell FPGA is used to interface the DCM board with the P0 dataway. The token passing method is used to pass data in the dataway. This allows output data to get built in the right order with the minimum overhead. The DSP host bus is used to connect to the FPGA. Because of the bandwidth sharing between link ports and host bus, the maximum bandwidth will be somewhere around 20 Megawords per second, if the input link port is running at full speed, with 40 Megawords per second burst speed.

F. DCM Crate

The DCM boards have a 6U X 280 mm VME form factor. The P0 connector in the VME64X backplane is used as a token passing 32 bits dataway. Eight signals including Strobe (clock), Valid, Token and Hold are used in controlling data from one DCM to the next DCM. The dataway is design to run at 40 MHz. The token

passing scheme enables the events to be built in the right sequence and at high bandwidth. Similar structure is also constructed in the P2 user defined space for passing LVL1 trigger primitives. Due to the lack of power provided to the VME backplane, 7 additional power pins are placed in P2 user defined space. Busies are also done in the similar fashion. A partition board is used to interface the crate to the Event Builder and LVL1 trigger system. This board determines how the DCM's get the partition in terms of data flow. Figure 4 shows a block diagram of the DCM crate.

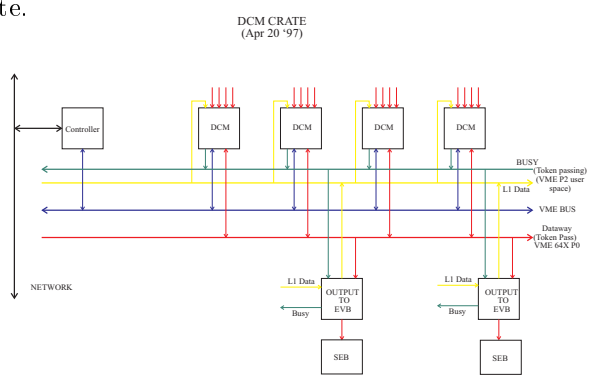

Fig. 4 Block diagram of DCM crate.

III. Recent Progress

The DCM board and two different compressor daughter boards have been fabricated. The AMU type compressor board and main board have been verified except for the data output through the P0 connector. A modified VME64X backplane will be ordered soon.

IV. References

[1] PHENIX Conceptual Design Report
[2] Analog Devices SHARC User Manual
[3] ALTERA FPGA Databook